

# Himanshu Bhardwaj

AI Engineer / ~3 years shipping production AI

himanshubhardwaj2398@gmail.com · LinkedIn · GitHub · himanshu-bhardwaj.com · +91-7053199308

AI engineer with ~3 years shipping production AI systems that domain experts can audit, not just trust. Most recently led AI extraction at **Bridge Medical** (Red Nucleus); previously **KalqylAI** — **4.6M** annual readers, **8,000+** AI equity reports per year.

Open to senior IC roles in production AI, and one freelance engagement per quarter.

## EXPERIENCE

### Bridge Medical Consulting (a Red Nucleus company)

Dec 2025 – Jun 2026

AI Engineer

Remote

- Own the AI technical strategy for HEOR Systematic Literature Review extraction — ~100 interdependent fields per study across population, intervention, comparator, outcome, design.
- **Track 1 (ship today):** containerize the existing notebook-based pipeline, cloud-deploy for non-technical reviewers, wire tracing, structured logging, and an evaluation/review surface.
- **Track 2 (design tomorrow):** anchor-based agentic pipeline — smart ingestion, validation, agentic escalation, citation-level provenance in a Postgres tree-and-graph schema. Every extracted value traces back to source page, sentence, or table cell.

### Predli Consulting AB

Sep 2023 – Dec 2025

AI & ML Consultant

Stockholm (Remote)

*KalqylAI (Spotlight Group) — autonomous AI equity research*

- Architected the event-driven Azure data layer: WebSocket ingestion for **160+** Swedish-listed companies, multi-format document parsing (PR + PDF + tables), Naive Bayes relevancy/PR classifier — **<30s** from press release to consumable event.
- Built the @KalqylAI Twitter agent: PR → published tweet in under 30 seconds, with entity-grounding (ticker, company, and subsidiary references reconciled against the structured data-layer payload before publish).
- Led the long-form analytical bot: dynamic outlines, rendered charts, financial ratios, segment commentary. **8,000+** AI equity reports, **4.6M** annual views, 16,000+ investors reached.

*Predli Studio — enterprise RAG platform*

- Built the document ingestion on Dagster: asset-based orchestration over a task scheduler, with content-hash idempotency at the raw-document boundary so connector retries don't poison the vector store.
- Designed the Enterprise RAG layer behind Studio: file-type-aware chunking (PDFs, decks, spreadsheets), TOC-based hierarchical context, persona-aware retrieval — **+40%** relevance over the uniform-chunking baseline.

### Predli Consulting AB

Apr 2023 – Jun 2023

Data Science Intern

Stockholm (Remote)

- End-to-end NLP pipeline for open-ended survey responses: Huggingface transformer embeddings → UMAP → HDBSCAN clustering. Themes emerge from the data, not predefined categories.

### ZS Associates

Nov 2019 – Jan 2022

Business Operations Associate

Gurgaon, India

- Designed incentive compensation frameworks for sales teams managing **\$100M+** annual revenue.
- Automated weekly performance reporting and COVID-19 impact dashboards — reduced manual effort by **20+ hrs/week**.

## EDUCATION

### Plaksha University

Technology Leaders Program (UC Berkeley & Purdue partnership)

2022 – 2023

CGPA 9.1/10

### Shiv Nadar University

B.Tech, Electrical & Electronics Engineering (Minor: Physics)

2015 – 2019

CGPA 8.08/10

## SELECTED STACK

**AI/ML:** LLMs, agentic AI, RAG, NLP, production ML, evaluation

**Languages & Data:** Python, SQL, PostgreSQL, vector databases

**Cloud & Infra:** Azure, AWS, Docker, CI/CD, Dagster, FastAPI

**Patterns:** event-driven ingestion, document parsing, structured extraction, audit trails